

Parallel in Time Simulation of Multiscale Stochastic Chemical Kinetics

Stefan Engblom¹ *

August 27, 2008

¹*Div of Scientific Computing, Dept of Information Technology
Uppsala University, P. O. Box 337, SE-75105 Uppsala, Sweden
email: stefane@it.uu.se*

Abstract

A version of the time-parallel algorithm parareal is analyzed and applied to stochastic models in chemical kinetics. A fast predictor at the macroscopic scale (evaluated in serial) is available in the form of the usual reaction rate equations. A stochastic simulation algorithm is used to obtain an exact realization of the process at the mesoscopic scale (in parallel).

The underlying stochastic description is a jump process driven by the Poisson measure. A convergence result in this arguably difficult setting is established suggesting that a homogenization of the solution is advantageous. We devise a simple but highly general such technique.

Three numerical experiments on models representative to the field of computational systems biology illustrate the method. For non-stiff problems, it is shown that the method is able to quickly converge even when stochastic effects are present. For stiff problems we are instead able to obtain fast convergence to a homogenized solution.

Overall, the method builds an attractive bridge between on the one hand, macroscopic deterministic scales and, on the other hand, mesoscopic stochastic ones. This construction is clearly possible to apply also to stochastic models within other fields.

Keywords: Parareal, reaction rate equations, jump process, homogenization, next reaction method, stochastic reaction-diffusion.

AMS subject classification: 65C40, 60J75, 60J22, 60H35, 68W10.

*Financial support has been obtained from the Swedish National Graduate School in Mathematics and Computing and the Swedish Foundation for Strategic Research.

1 Introduction

It has been shown in several studies [13, 36, 38, 40] that stochastic descriptions of biochemical networks are necessary for understanding and explaining the mechanisms inside living cells. Such networks often contain species present in copy numbers down to a few hundreds [25], hinting why discreteness and stochasticity becomes important. Examples of when random effects are pronounced include the regularity and drift of circadian oscillations [46], spontaneous separation in bistable systems [12] and the creation of new steady-states [45].

An accurate such stochastic model can be obtained directly from microphysical considerations as a consequence of the Markov property [23]. Essentially, in a diffusion-controlled system, reactive collisions between molecules are rare events, implying a rapid loss of dependence on the systems history.

Sample realizations of any given chemical network can be obtained by randomly firing reactions and correspondingly updating the systems state; the most well-known such algorithm is Gillespie's *stochastic simulation algorithm* (SSA) [22]. Although conceptually simple, detailed simulation of complex reaction networks as found for example in living cells remains a computationally very intensive problem.

The *parareal* algorithm for the solution of evolution equations on a time-interval $[0, T]$ was suggested in the note [33] as a method for the *parallel* solution of problems in *real* time. An improved version, better tuned to nonlinear problems, was subsequently applied to problems in control theory [35] and molecular dynamics [3] to mention just a few.

The method is built around a predictor-corrector step in which a *coarse* solver is used as a preconditioner to a *fine* solver. It thus incorporates both multigrid and domain decomposition ideas, but is unique in that it is a purely parallel algorithm; it has no value when executed on a serial machine.

The idea suggested in this paper is to use the macroscopic approximation of the chemical system, the *reaction rate equations*, as the coarse solver while a stochastic simulation is used as the fine scale solver. This setup is available also for other types of stochastic evolution equations but has not received much attention. The reason is probably that any deterministic model is a quite poor formal approximation. On the other hand, any macroscopic trend in the solution is likely to be quickly obtained. Also, with a deterministic coarse solver, highly efficient implicit integrators are available which greatly simplifies the implementation.

The present paper is related to references [5, 34]. In the former, the analysis of parareal when applied to ordinary differential equations (ODEs) and stochastic ODEs (SDEs) driven by Wiener processes is treated. In the lat-

ter, the parareal algorithm is applied to linear mass-balance (deterministic) chemical kinetics with disparate rates by using a coarse solver in the form of a reduced model.

The paper is divided into sections as follows: in Section 2 the mathematical model of stochastic kinetics is described in some detail and we briefly review the parareal algorithm in Section 3. We present a convergence analysis in Section 4; the main result indicates that strong convergence might be slow for systems with large Lipschitz constants. A way around this is suggested in the form of a simple but very general homogenization procedure in parallel. This “parareal homogenization” could well be an interesting approach to many other stiff problems. We perform three quite different numerical experiments in Section 5 where the practical implementation of the algorithm is also discussed. Finally, in Section 6 we conclude the paper by discussing the outcome of the experiments and suggesting some future work and generalizations.

2 Stochastic chemical kinetics

In this section we discuss the usual mesoscopic stochastic model for chemical reactions (the master equation). As it is needed in the analysis in Section 4, we shall also give an equivalent and mathematically more appealing representation in terms of a certain stochastic jump process. For monographs on applications of the master equation, see [18, 29]. The mathematical treatment of jump processes can be found in [2, 9, 15, 21, 27].

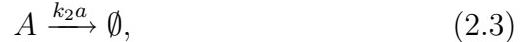
2.1 Chemical reactions, the master equation and a hierarchy of methods

We consider in this paper a general homogenous chemical network consisting of D different species reacting according to R prescribed reactions. At any given time t , the *state* of the system is an integer vector $x \in \mathbf{Z}_+^D = \{0, 1, 2, \dots\}^D$ counting the number of molecules of each kind. In a stochastic description, each reaction is a change of the state according to a certain transition rule and the intensity of this process is governed by a *reaction propensity*, $w_r : \mathbf{Z}_+^D \rightarrow \mathbf{R}_+$. This is the transition probability per unit of time for moving from the state x to $x - \mathbb{N}_r$,

$$x \xrightarrow{w_r(x)} x - \mathbb{N}_r, \quad (2.1)$$

where by convention, $\mathbb{N}_r \in \mathbf{Z}^D$ is the transition step and is the r th column in the *stoichiometric matrix* $\mathbb{N} \in \mathbf{Z}^{D \times R}$.

For instance, in a given well-stirred volume Ω the *elementary* reactions can be written using the states $x = [a, b]^T$,



The propensities are generally scaled such that

$$w_r(x) = \Omega u_r(x/\Omega) \quad (2.6)$$

for some function u_r which does not depend on Ω . Intensities of this form are called *density dependent* and arise naturally in a number of situations [15, Ch. 11].

Arguably the most common stochastic description of chemically reacting systems is the *chemical master equation* (CME). To state it, let $p(x, t)$ be the probability that a certain number x of molecules is present at time t . The CME [18, 29] is then given by

$$\frac{\partial p(x, t)}{\partial t} = \sum_{\substack{r=1 \\ x+\mathbb{N}_r^- \geq 0}}^R w_r(x + \mathbb{N}_r)p(x + \mathbb{N}_r, t) - \sum_{\substack{r=1 \\ x-\mathbb{N}_r^+ \geq 0}}^R w_r(x)p(x, t), \quad (2.7)$$

where the transition steps are decomposed into positive and negative parts as $\mathbb{N}_r = \mathbb{N}_r^+ + \mathbb{N}_r^-$.

In a pioneering paper from 1976, Gillespie [22] showed that exact samples from the master equation can be obtained via the SSA. This follows essentially by identifying it with the forward Kolmogorov equation associated to a certain continuous-time Markov chain to be introduced shortly. Importantly, Gillespie [23] later also showed that the CME is an exact physical description when the system is well-stirred and in thermal equilibrium.

For chemical networks where the number of reactions per unit of time is large, detailed simulation with SSA becomes inefficient. As a remedy, Gillespie [24] proposed the *tau-leap method* with the ability to simultaneously “leap” over many reactions at once. The method has since been developed and improved by several authors [1, 31, 42].

A related issue is *stiffness*: many interesting models become prohibitively expensive to solve by explicitly simulating the various involved scales. Several different model reduction techniques have been proposed for this situation [7,

11, 26]. Also, an implicit version of the tau-leap method has been developed [41], but this method converges in a very weak sense only [8, 32, 42].

An alternative approximation is the *Langevin equation* [15, Ch. 7] which is a *continuous* SDE driven by Wiener processes. On ignoring the noise term and only keeping the drift, the *reaction rate* equations are obtained. This is a set of D ODEs approximately evolving the expectation value of the mesoscopic model. The hierarchy of solution methods (SSA → tau-leap → Langevin SDE → ODE) can thus be thought of as a transition between the mesoscopic model and the macroscopic one [24].

2.2 Sample path representation

A representation equivalent to, but perhaps more direct than the master equation (2.7), was proposed fairly recently in [39]. The idea is to construct a sample path representation in the form of jump SDEs driven by Poisson processes. As we shall see and as pointed out in [31, 39], this representation is particular useful for numerical analysis.

We thus introduce the stochastic variable $X(t) \in \mathbf{Z}_+^D$ counting at time $t \geq 0$ the number of molecules of each type. Reactions are understood to occur instantly and thus the process is right continuous only; the notation $X(t-)$ is used to denote the value of the process prior to any reactions occurring at time t . We assume the existence of a probability space¹ $(\Sigma, \mathbf{F}, \mathbf{P})$ with the filtration $\mathbf{F}_{t \geq 0}$ containing R -dimensional standard Poisson processes. The transition probabilities in (2.1) define *counting processes* $\pi_r(t)$ [9, Ch. 2.5] according to

$$E[\pi_r(t + dt) - \pi_r(t)|\mathbf{F}_t] = w_r(X(t-)) dt + o(dt). \quad (2.8)$$

In turn, the counting processes determine the process $X(t)$ [15, Ch. 6.4],

$$X_t = X_0 - \sum_{r=1}^R \mathbb{N}_r \pi_r(t). \quad (2.9)$$

A representation of $\pi_r(t)$ in terms of a unit-rate Poisson process Π_r in an *operational* or *scaled* time can also be given [9, Ch. 3.1, 4.4],

$$\pi_r(t) = \Pi_r \left(\int_0^t w_r(X(s-)) ds \right). \quad (2.10)$$

¹Note that, throughout the paper, the letter Ω is *not* used for the probability sample space, but consistently denotes the system volume.

We now wish to move beyond (2.9) and obtain a more transparent representation in the form of a jump SDE driven by the proper measure. Processes with nonlinearly dependent jump intensities are quite difficult to handle since the noise enters with a complicated dependence on the state. This is in contrast to the more familiar Itô SDE driven linearly by Wiener processes.

For the present purposes, the Poisson random measure [28, Ch. II.1c] $\mu(dt \times dz; \sigma)$ with $\sigma \in \Sigma$ defines an increasing sequence of arrival times $\tau_i \in \mathbf{R}_+$ with corresponding “marks” z_i uniformly distributed in $[0, \bar{W}]$ [9, Ch. 1.4], where \bar{W} is to be defined shortly. The deterministic intensity of $\mu(dt \times dz)$ is the Lebesgue measure, $m(dt \times dz) = dt \times dz$ [21, Part II, Ch. 2§4].

The propensities have been general nonlinear functions up to this moment but we need to impose conditions so as to bound the total intensity of all reaction channels. It is not difficult to see that a *finite* intensity implies a bounded solution in the mean square sense for finite times (see Lemma 4.2 below). For convenience, we take the approach in [31] and formally specializes the investigation to closed systems. This implies the existence of a bounded set S containing at any time t the state of the system and is reasonable from physical considerations; —any real application must necessarily involve a bounded total number of molecules.

Summarizing, we thus have

Assumption 2.1. The state of the chemical network satisfies $X(t) \in S \subset \mathbf{Z}_+^D$. When confined to this set, all propensities are Lipschitz continuous in their argument with respect to the Euclidean norm, $|w_r(x) - w_r(y)| \leq L_r \|x - y\|$. It follows that all propensities are bounded such that

$$\bar{W}_r := \sum_{s=1}^r \max_{x \in S} w_s(x) \quad (2.11)$$

exists and is bounded for all r . We define also the numbers $L = \sum_r L_r$, $\bar{W} = \bar{W}_R$ and $W = \max_{x \in S} \sum_r w_r(x)$. Note that $W \leq \bar{W}$.

The frequency of each reaction is controlled through a set of indicator functions $\hat{w}_r : S \times [0, \bar{W}] \rightarrow \{0, 1\}$ defined as follows:

$$\hat{w}_r(x; z) = \begin{cases} 1 & \text{if } 0 \leq z - \bar{W}_{r-1} < w_r(x), \\ 0 & \text{otherwise.} \end{cases} \quad (2.12)$$

We can now represent the counting process (2.10) in terms of the Poisson random measure [28, Ch. II.1d] via a *thinning* of the measure using an acceptance-rejection strategy [9, Ch. 4.3]:

$$\pi_r(t) = \int_0^t \int_0^{\bar{W}} \hat{w}_r(X(t-); z) \mu(dt \times dz). \quad (2.13)$$

Note that this thinning is not the same as in [31, 39] where instead $z \in [0, W]$; the present construction yields sharper estimates when comparing trajectories formed from different initial data.

The representation (2.13) combined with (2.9) leads us to the sample path representation in terms of a Skorohod jump SDE [27, Ch. IV.9],

$$dX_t = - \sum_{r=1}^R \mathbb{N}_r \int_0^{\bar{W}} \hat{w}_r(X(t-); z) \mu(dt \times dz). \quad (2.14)$$

Often one is interested in separating (2.14) into its “drift” and “jump” terms [21, Part II, Ch. 2§4],

$$\begin{aligned} dX_t &= - \sum_{r=1}^R \mathbb{N}_r w_r(X(t-)) dt - \sum_{r=1}^R \mathbb{N}_r \int_0^{\bar{W}} \hat{w}_r(X(t-); z)(\mu - m)(dt \times dz) \\ &=: dX_D(t) + dX_J(t). \end{aligned} \quad (2.15)$$

We will use the initial value convention $X_D(0) = X(0)$ and $X_J(0) = 0$ throughout the paper. Note that the second term in (2.15) driven by the compensated measure $(\mu - m)$ is a martingale of mean zero. On taking expectation values and approximating, we obtain $dEX_t \approx d\tilde{x}(t)$ where

$$d\tilde{x}(t) = - \sum_{r=1}^R \mathbb{N}_r w_r(\tilde{x}(t)) dt \quad (2.16)$$

and $\tilde{x}(0) = EX(0)$. This approximation constitutes, up to a scaling by the system volume Ω , the usual macroscopic reaction-rate equations.

3 The parareal algorithm

We continue by giving a short account of the parareal algorithm. For a quick introduction with additional references the survey [43] can be recommended. See also [4, 17, 44] for abstract convergence results, stability analysis and more.

Consider the general time-dependent problem written in operator form,

$$\dot{u} = -Au, t \in [0, T] \text{ with } u(0) = u_0. \quad (3.1)$$

Write the solution propagator \mathcal{F} as

$$\mathcal{F}_t(y) = y - \int_0^t Au(t) dt, \quad (3.2)$$

where u solves (3.1) subject to $u(0) = y$. The parareal algorithm assumes that a coarse solver \mathcal{C} is available that does the same thing but faster and presumably less accurate.

Time is now discretized in $N = T/\Delta t$ chunks of time and any solver $S \in \{\mathcal{F}_{\Delta t}, \mathcal{C}_{\Delta t}\}$ can be used to compute a numerical solution:

$$\begin{bmatrix} I & 0 & 0 & 0 \\ -S & I & 0 & 0 \\ 0 & -S & I & 0 \\ 0 & 0 & -S & I \end{bmatrix} \begin{bmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} u_0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad (3.3)$$

or simply $B(S)v = u_0$ in matrix notation. The parareal algorithm emerges as the fix-point iteration obtained by using $B(\mathcal{C}_{\Delta t})^{-1}$ as an approximate inverse to $B(\mathcal{F}_{\Delta t})$:

$$v_{k+1} = v_k - B(\mathcal{C}_{\Delta t})^{-1}(B(\mathcal{F}_{\Delta t})v_k - u_0). \quad (3.4)$$

Let $v_{0,0} = u_0$ and $v_{0,n} = \mathcal{C}_{\Delta t}v_{0,n-1}$ for $n \in [1, N]$ to start up the algorithm. One readily verifies the recursion

$$v_{k,n} = \mathcal{F}_{\Delta t}v_{k-1,n-1} - [\mathcal{C}_{\Delta t}v_{k-1,n-1} - \mathcal{C}_{\Delta t}v_{k,n-1}], \quad (3.5)$$

where the evaluation of \mathcal{F} is trivially parallel.

At any point in the algorithm, the preconditioned residual is given by

$$p_k = v_k - v_{k+1}, \quad (3.6)$$

and is useful as an estimate of the error.

A nice survey of the parareal method is found in [17] where the connection to some earlier methods is made. It turns out that one can view the process as a Newton method applied to a certain nonlinear problem using a Jacobian approximated by finite differences. This partially explains the superlinear convergence often observed on bounded intervals. It is known that the parareal algorithm converges poorly for problems where the operator has eigenvalues with large imaginary parts [44]. Another interesting result is that the parareal scheme can be made unconditionally stable by using a suitable relaxation of the coarse solver [4].

4 Analysis

In this section we analyze the parareal algorithm obtained by using the rate equations (2.16) and the jump SDE (2.14) as the coarse and fine solver,

respectively. In Section 4.1 we give some results that cover local properties of solutions to these equations. The actual convergence analysis is found in Section 4.2 where convergence in mean square and convergence of the first moment is investigated. As noted in [5], weak convergence is not so interesting for the parareal algorithm since weak estimates typically come into play when many trajectories are simultaneously computed in a Monte Carlo simulation; parallelism can then be trivially achieved with optimal efficiency. Our interest in the convergence of the first moment mainly stems from the possibility to homogenize the model on the fly by a suitable averaging filter. This idea is discussed in Section 4.3.

Because the process $X(t)$ is integer valued, we mention here for clarity the trivial inequality

$$|X| \leq X^2 \implies E|X| \leq EX^2. \quad (4.1)$$

When X is large, (4.1) is an overestimate and we will switch to the standard inequality

$$E|X| \leq (EX^2)^{1/2}. \quad (4.2)$$

For simplicity and without serious loss of generality, we treat the 1-D case only.

4.1 Preliminaries

We start by citing the following convenient lemma.

Lemma 4.1 (Lemma 3.1 and 3.2 in [31]; see also Lemma 2.3.2 in [2]). *Define $\Delta X_t = X(t) - X(t-)$. Then for any fixed time $s > 0$, $\Delta X_s = 0$ (a.s.). Furthermore, if w is a continuous function and $t_1 < t_2$, then*

$$\int_{t_1}^{t_2} \Delta w(X_t) dt = 0. \quad (4.3)$$

We continue by establishing two lemmas concerning the stability of the jump process $X(t)$.

Lemma 4.2. *Let $X(t \geq 0)$ evolve according to (2.14) and define the jump term $X_J(t)$ as in (2.15) with $X_J(0) = 0$. Then*

$$EX_J(t)^2 \leq \|\mathbb{N}\|^2 Wt. \quad (4.4)$$

Furthermore,

$$E[X_t - X_0]^2 \leq 2\|\mathbb{N}\|^2 Wt + 2\|\mathbb{N}\|^2 W^2 t^2. \quad (4.5)$$

Proof. The left side of (4.4) can be written

$$\begin{aligned} & E \left(\int_0^t \int_0^{\bar{W}} \sum_{r=1}^R -\mathbb{N}_r \hat{w}_r(X(s-); z)(\mu - m)(ds \times dz) \right)^2 \\ &= E \int_0^t \int_0^{\bar{W}} \left(\sum_{r=1}^R -\mathbb{N}_r \hat{w}_r(X(s-); z) \right)^2 m(ds \times dz) \\ &\leq E \int_0^t \int_0^{\bar{W}} \|\mathbb{N}\|^2 \sum_{r=1}^R \hat{w}_r(X(s-); z)^2 m(ds \times dz), \end{aligned}$$

where the Itô isometry for jump processes was used [21, Part II, Ch. 2§5] [2, Lemma 4.2.2]. Using that $\hat{w}_r^2 = \hat{w}_r$ we complete the proof of (4.4). For (4.5) we use the drift/jump split (2.15) and the inequality $(a+b)^2 \leq 2a^2 + 2b^2$ to obtain

$$E[X_t - X_0]^2 \leq 2EX_J(t)^2 + 2E \left(\int_0^t \sum_{r=1}^R -\mathbb{N}_r w_r(X(s-)) ds \right)^2.$$

□

Lemma 4.3. *Let the drift term be defined as in (2.15) with $X_D(0) = X(0)$ and let $\tilde{x}(t)$ be the reaction rate approximation (2.16) to the expected value of $X(t)$ with $\tilde{x}(0) = X(0)$. Then*

$$E[X_D(t) - \tilde{x}(t)]^2 \leq L^2 \|\mathbb{N}\|^4 W t^3 \exp(2L^2 \|\mathbb{N}\|^2 t^2). \quad (4.6)$$

Proof. By Jensen's inequality and the Lipschitz boundedness we get

$$\begin{aligned} E[X_D(t) - \tilde{x}(t)]^2 &\leq L^2 \|\mathbb{N}\|^2 t E \int_0^t [X(s) - \tilde{x}(s)]^2 ds \\ &\leq L^2 \|\mathbb{N}\|^2 t \int_0^t 2EX_J(s)^2 + 2E[X_D(s) - \tilde{x}(s)]^2 ds \\ &\leq L^2 \|\mathbb{N}\|^4 W t^3 + 2L^2 \|\mathbb{N}\|^2 t \int_0^t E[X_D(s) - \tilde{x}(s)]^2 ds, \end{aligned}$$

where the first part of Lemma 4.2 was used. The result now follows as an application of Grönwall's inequality. □

Remark. In a closed system we can identify the magnitude of X with the system volume Ω ; $\max_{x \in S} \|x\| = \mathcal{O}(\Omega)$. For density dependent propensities (see (2.6)), this implies that $W \sim \Omega$ while L does not scale with Ω . Lemma 4.2 and 4.3 then yields

$$(E[X(t) - \tilde{x}(t)]^2)^{1/2} \leq C_t \Omega^{1/2} \quad (4.7)$$

for some bounded constant C_t . By Chebyshev's inequality we thus have convergence in probability $\Omega^{-1}X(t) \rightarrow \Omega^{-1}\tilde{x}(t)$ as $\Omega \rightarrow \infty$. For more precise reasoning of this type, see [15].

We proceed with a result that is crucial to the convergence properties of the parareal algorithm.

Theorem 4.4. *Let $X(t \geq 0)$ and $Y(t \geq 0)$ evolve according to (2.14) using the same underlying set of Poisson processes but with different initial values X_0 and Y_0 . Define the jump terms $X_J(t)$ and $Y_J(t)$ as in Lemma 4.2. Then*

$$E[Y_J(\Delta t) - X_J(\Delta t)]^2 \leq L\|\mathbb{N}\|^2\Delta t|Y_0 - X_0|(1 + \mathcal{O}(\Delta t)). \quad (4.8)$$

Proof. By the Itô isometry we can bound the left side of (4.8) by

$$\begin{aligned} & E \int_0^{\Delta t} \int_0^{\bar{W}} \|\mathbb{N}\|^2 \sum_{r=1}^R [\hat{w}_r(Y(t-); z) - \hat{w}_r(X(t-); z)]^2 m(dt \times dz) \\ & \leq E \|\mathbb{N}\|^2 \int_0^{\Delta t} \sum_{r=1}^R |w_r(Y(t-)) - w_r(X(t-))| dt, \end{aligned}$$

where the usefulness of the thinning proposed in Section 2.2 is evident. From the Lipschitz assumption and Lemma 4.1 we get the bound

$$\begin{aligned} & \leq L\|\mathbb{N}\|^2 \int_0^{\Delta t} E|Y(t-) - X(t-)| dt = L\|\mathbb{N}\|^2 \int_0^{\Delta t} E|Y(t) - X(t)| dt \\ & \leq L\|\mathbb{N}\|^2 \Delta t |Y_0 - X_0| + L\|\mathbb{N}\|^2 \int_0^{\Delta t} E|Y(t) - Y_0| + E|X(t) - X_0| dt \\ & \leq L\|\mathbb{N}\|^2 \Delta t |Y_0 - X_0| + 4L\|\mathbb{N}\|^4 \int_0^{\Delta t} Wt + W^2t^2 dt, \end{aligned}$$

where the integer inequality (4.1) and the second part of Lemma 4.2 was used in the last line. \square

The following result, similar in spirit to Theorem 4.4, is useful for studying the error in the first moment alone.

Theorem 4.5. *Define $Y_D(t)$, $\tilde{y}(t)$, $X_D(t)$ and $\tilde{x}(t)$ as in Lemma 4.3. Then*

$$|EY_D(\Delta t) - \tilde{y}(\Delta t) - EX_D(\Delta t) + \tilde{x}(\Delta t)| \leq \mathcal{O}(\Delta t^{3/2}) |E[Y_0 - X_0]|. \quad (4.9)$$

Proof. If $Y_0 = X_0$ there is nothing to prove and we may thus assume that $|Y_0 - X_0| \geq 1$. Two applications of Lemma 4.3 gives that

$$Y_D(\Delta t) - \tilde{y}(\Delta t) - X_D(\Delta t) + \tilde{x}(\Delta t) = \mathcal{O}(\Delta t^{3/2}).$$

We recover (4.9) by inserting the factor $(Y_0 - X_0)$, taking expectation and absolute values of both sides. \square

4.2 Convergence of the parareal algorithm

For the application of the parareal algorithm we let $\mathcal{F}y$ be the stochastic evolution in a time-step Δt of the jump process $X(t)$ obeying (2.14) with initial data y . Following (2.15), we write $\mathcal{F} = \mathcal{F}_D + \mathcal{F}_J$ for the drift- and jump term respectively. The coarse solver $\mathcal{C}y$ is instead the evolution in a time-step Δt using the deterministic rate equations (2.16) starting from y .

We denote by $X = [X_0, \dots, X_N]$ with $X_n = X(t_n)$ the exact solution at times $t_n \equiv n\Delta t$. $\tilde{X}_{k,n} \approx X_n$ is the numerical approximative solution obtained after the k th iteration of the parareal algorithm:

$$\tilde{X}_{k,0} = X_0, \quad k \geq 0, \quad (4.10)$$

$$\tilde{X}_{0,n} = \mathcal{C}\tilde{X}_{0,n-1}, \quad n \geq 1, \quad (4.11)$$

$$\tilde{X}_{k,n} = \mathcal{F}\tilde{X}_{k-1,n-1} - \mathcal{C}\tilde{X}_{k-1,n-1} + \mathcal{C}\tilde{X}_{k,n-1}, \quad (4.12)$$

where $(k, n) \geq 1$ in (4.12).

We now wish to analyze the root-mean-square (RMS) error defined by

$$e_{k,n}^2 = E[\tilde{X}_{k,n} - X_n]^2. \quad (4.13)$$

This measure satisfies by (4.12) the recursion

$$\begin{aligned} e_{k,n}^2 &= E[T_1 + T_2 + T_3]^2 \\ &= E[T_1^2 + T_2^2 + T_3^2 + 2T_1T_2 + 2T_1T_3 + 2T_2T_3], \end{aligned} \quad (4.14)$$

where in terms of

$$T_1 \equiv \mathcal{F}_J\tilde{X}_{k-1,n-1} - \mathcal{F}_JX_{n-1}, \quad (4.15)$$

$$T_2 \equiv \mathcal{F}_D\tilde{X}_{k-1,n-1} - \mathcal{C}\tilde{X}_{k-1,n-1} - \mathcal{F}_DX_{n-1} + \mathcal{C}X_{n-1}, \quad (4.16)$$

$$T_3 \equiv \mathcal{C}\tilde{X}_{k,n-1} - \mathcal{C}X_{n-1}. \quad (4.17)$$

By Theorem 4.4 using (4.1) and Lemma 4.3 we get

$$ET_1^2 \leq L\|\mathbb{N}\|^2\Delta te_{k-1,n-1}^2(1 + \mathcal{O}(\Delta t)), \quad (4.18)$$

$$ET_2^2 \leq \mathcal{O}(\Delta t^3). \quad (4.19)$$

We also have the Lipschitz bound

$$ET_3^2 \leq \exp(2L\|\mathbb{N}\|\Delta t)e_{k,n-1}^2. \quad (4.20)$$

As for the cross-terms we have that $ET_1T_3 = 0$ since T_1 is a martingale of zero mean and T_3 is a deterministic (non-anticipating) function. The remaining terms can be easily estimated as

$$2ET_1T_2 \leq \Delta t ET_1^2 + \Delta t^{-1} ET_2^2 \leq \mathcal{O}(\Delta t^2) e_{k,n-1}^2, \quad (4.21)$$

$$2ET_2T_3 \leq \Delta t^{-1} ET_2^2 + \Delta t ET_3^2 \leq \mathcal{O}(\Delta t^2) + \mathcal{O}(\Delta t) e_{k,n-1}^2. \quad (4.22)$$

Summarizing, we have from (4.14) upon ignoring higher order terms in Δt ,

$$\begin{aligned} e_{k,n}^2 &\leq L\|\mathbb{N}\|^2 \Delta t e_{k-1,n-1}^2 + \exp(2L\|\mathbb{N}\|\Delta t) e_{k,n-1}^2 \\ &=: \Delta t S_{\mathcal{F}}^2 e_{k-1,n-1}^2 + \exp(2\Delta t S_{\mathcal{C}}) e_{k,n-1}^2. \end{aligned} \quad (4.23)$$

Define $M = \max_n e_{0,n}$ in view of (4.7). We readily recognize the binomial recurrence in (4.23) so that

$$e_{k,n}^2 \leq M^2 \binom{n}{k} \Delta t^k S_{\mathcal{F}}^{2k} \exp(2\Delta t S_{\mathcal{C}}(n-k)). \quad (4.24)$$

If we now look at iteration $k = \mathcal{O}(1)$ and interval $n = N = T/\Delta t$, (4.24) can be simplified into

$$e_{k,n} \leq C_{1,T} S_{\mathcal{F}}^k, \quad (4.25)$$

where $C_{1,T}$ is a bounded constant for any given total time T . Eq. (4.25) shows mean square convergence for contractive problems where $S_{\mathcal{F}} < 1$ only; it is unclear what happens for systems with larger Lipschitz constants. A key to understanding how the analysis can be refined lies in the fact that the integer inequality (4.1) had to be applied before using Theorem 4.4 and arriving at (4.18). For an initial large error, the inequality (4.2) is sharper but introduces a nonlinear dependence:

$$e_{k,n}^2 \leq \Delta t S_{\mathcal{F}}^2 e_{k-1,n-1} + \exp(2\Delta t S_{\mathcal{C}}) e_{k,n-1}^2. \quad (4.26)$$

This motivates our interest in the next proposition.

Proposition 4.6. *Consider for $(k, n) \geq 0$ a sequence $a_{k,n}$ of nonnegative numbers. Let $a_{k,0} = 0$ for $k \geq 0$ and suppose that $a_{0,n} \leq C$ for $n > 0$ and some nonnegative constant C . Then for $k \leq n$ the inequality*

$$a_{k,n} \leq A a_{k-1,n-1}^{1/2} + B a_{k,n-1}$$

where $A \geq 0$ and $B \geq 1$ is satisfied by

$$a_{k,n} \leq A^{2^{2-2^{1-k}}} B^{n-k} (n-k)^{2^{2-2^{1-k}}} C^{2^{-k}}. \quad (4.27)$$

A proof by induction is easily constructed once the functional form in (4.27) has been obtained.

To apply Proposition 4.6 we put $A = \Delta t S_{\mathcal{F}}^2$, $B = \exp(2\Delta t S_{\mathcal{C}})$ and identify $e_{k,n} = a_{k,n}^{1/2}$ so that $C = M^2$. The result is for $k \ll n = N$ that

$$e_{k,n} \leq C_{2,T} M^{2^{-k}}, \quad (4.28)$$

for some constant $C_{2,T}$. During the first few iterations, this bound is sharper than (4.25) which is only valid when the error is already small. For instance, if as in (4.7) we have that $M \sim \Omega^{1/2}$, then $e_{k,N} \sim [\Omega^{1/2}, \Omega^{1/4}, \Omega^{1/8}, \dots]$ for the first few iterations $k = [0, 1, 2, \dots]$.

As it provides us with additional insight, we also comment on the *error in the first moment*. Define for this purpose $\bar{e}_{k,n} = |E[\tilde{X}_{k,n} - X_n]|$. Proceeding as before and using Theorem 4.5 it is not difficult to see that this error satisfies

$$\begin{aligned} \bar{e}_{k,n} &\leq \mathcal{O}(\Delta t^{3/2}) \bar{e}_{k-1,n-1} + \exp(L\|\mathbb{N}\|\Delta t) \bar{e}_{k,n-1} \\ &\lesssim \binom{n}{k} \Delta t^{3k/2} \leq C_{3,T} \Delta t^{k/2}. \end{aligned} \quad (4.29)$$

The estimates (4.25), (4.28) and (4.29) show that *for Δt sufficiently small, the contractive parts converge in mean square while the non-contractive parts converge in the sense of (4.29) only*. This interpretation opens up for modifying the scheme in such a way that the fast scales are actively filtered away.

4.3 Homogenization

The previous analysis suggests that problems with large Lipschitz constants may be problematic to solve by the proposed parareal algorithm. On the other hand, for certain problems involving very rapid scales, pathwise convergence may not be so interesting to obtain. Rather are we interested in convergence to a *homogenized* model which by itself is a weak approximation to the original system. For example, this situation occurs in the current context when rapid reaction channels almost balance each other so that the interesting dynamics occurs on a much slower scale.

Unlike the deterministic case where implicit solvers generally evolve stiff problems efficiently, it has been proposed that model reduction techniques are necessary in the stochastic setting [32]. For instance, the implicit tau-leap method has only been shown to be weakly convergent under the rather strong assumption of linear propensities [8, 42].

Provided that the coarse solver is sufficiently dissipative, the biggest challenge in directly using the current parareal algorithm for stiff problems is to

detect when the numerical solution is of sufficient quality; since convergence is at best weak, any jumps on the order of the noise associated with the fast scale could in principle be accepted.

An easy way around this is to replace the fine propagator \mathcal{F} with a homogenized version \mathcal{F}^h . A simple example is

$$\mathcal{F}^h X_0 := \frac{1}{\delta t} \int_{\Delta t - \delta t}^{\Delta t} Y(t) dt, \quad \text{where } Y(t) = \mathcal{F}_t X_0, \quad (4.30)$$

i.e. an average of the exact trajectory over an interval δt . This interval should be large enough to contain several fast reactions but short enough to be essentially independent on the slow scales. One can easily think of much more advanced filters to create a suitable homogenization but we settled for the immediate (4.30).

Let us remark that an essentially “exact” trajectory is available within the homogenized solution since the non-averaged solution is always computed first. This “exact” trajectory is just as noisy as the true solution but allows for jumps related to the fast scale in between successive time intervals $[t_n, t_{n+1}]$. An intuitive interpretation is that the homogenized trajectory is an exact sample from a nearby model containing additional reactions that are scheduled at deterministic time-steps Δt . In principle, the intensity of this unknown process could be estimated and be put in relation to the rest of the propensities.

5 Numerical examples

After discussing the implementation of the proposed method, we will consider three numerical examples. The first is a typical example of when stochastic models are necessary since mean-field equations generally give incorrect results. Nevertheless, when they are used as a preconditioner in the parareal algorithm, convergence to the true solution is quite fast. The second example is representative of situations involving multiple scales with fast transients. Finally, we obtain a trajectory from the *reaction-diffusion* master equation which is a good representative of very large networks. Here, the macroscopic model is just the familiar reaction-diffusion partial differential equation.

5.1 Implementation

For the purpose of performing experiments we have implemented a serial version of the parareal algorithm in Matlab. The coarse solver is thus simply the reaction rate equations solved by any suitable ODE-solver; we typically

used Matlab's `ode23s` but occasionally got away with just a single step of the backward Euler method. As the fine solver we used an implementation of the *next reaction method* (NRM) [20], which is statistically equivalent to the SSA but specifically tuned to large networks. The reason for preferring this particular implementation is that *consistent Poisson processes may easily be sampled*. This feature is achieved by simply using the same sequence of random numbers for each reaction channel.

To appreciate where this technicality enters in the analysis, note that in Theorem 4.4, Y_t and X_t are assumed to be trajectories conditioned on the same Poisson process. If they are not related in this respect, the associated constants are much larger.

In order to estimate the quality of the numerical solution we used the relative Euclidean norm of the preconditioned residual as given by (3.6);

$$e_{k+1} \sim \max_n D^{-1/2} \|(v_{k,n} - v_{k+1,n})/(1 + v_{k+1,n})\|, \quad (5.1)$$

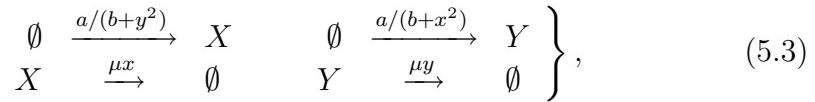
(elementwise division). This was then taken as an approximation to the true relative error

$$e_{k+1} = \max_n D^{-1/2} \|(v_{k+1,n} - u_n)/(1 + u_n)\|. \quad (5.2)$$

We generally did not round any fractional results obtained from the coarse solver. Formally, this introduces a complication in the analysis as the solution space becomes a continuum (see [31, Remark 3.1]) but we have not found any benefits in forcing the solution to stay in \mathbf{Z}_+^D . However, for fractional states, the propensities should explicitly be set to zero whenever executing the reaction would result in a negative copy number.

5.2 Stochastic toggle switch

A *toggle switch* found within the regulatory network of *E. coli* has been modeled by two mutually cooperatively repressing gene products X and Y [19]. The model is



with parameters $a = 3000$, $b = 11000$ and $\mu = 10^{-3}$. In (5.3), note that if the number of X -molecules is larger than the number of Y -molecules, then the production of Y -molecules is inhibited and the system finds a stable state with $x > y$. However, the intrinsic noise can make the roles of X and Y

suddenly switch to a state where instead the production of X -molecules is inhibited. The system (5.3) thus constitutes an interesting example where the deterministic dynamics clearly differs from that obtained by a stochastic simulation.

In Figure 5.1 the exact trajectory up to final time $T = 5 \times 10^6$ is displayed together with the approximation obtained after a few iterations by the parareal algorithm on a *very* coarse grid with $\Delta t = T/50 = 10^5$. Interestingly, the next correct place to switch is found for each new iteration so that all such events are correctly located within the first 4 iterations. Although the numerical solution sometimes overshoots, *some* information is evidently being correctly propagated through the system.

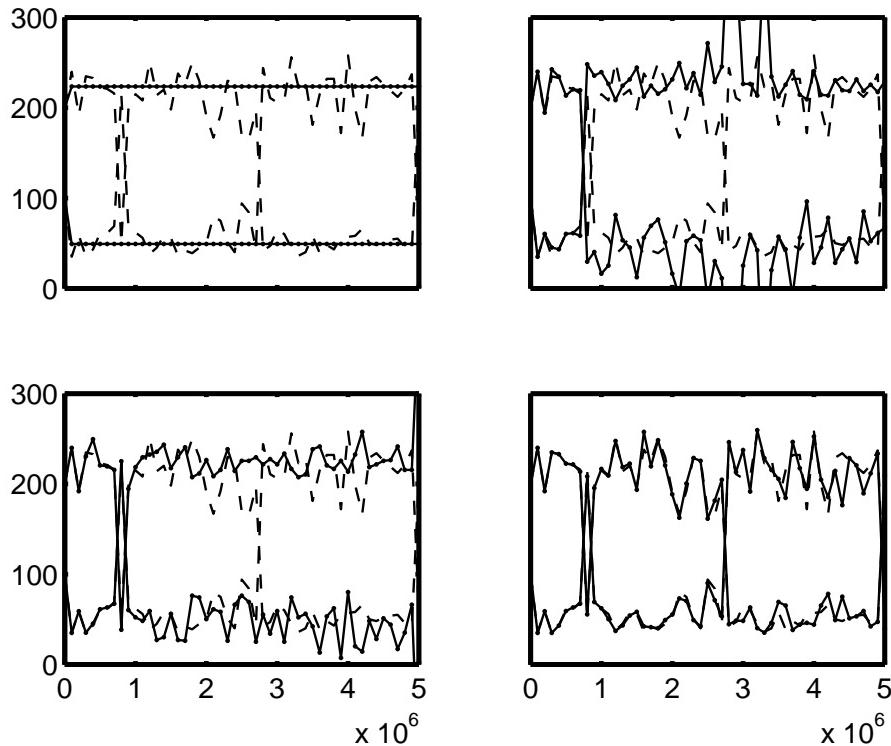


Figure 5.1: Dashed: exact trajectory of the two components of the toggle switch (5.3) versus time. Solid: solution obtained in parallel after 0, 1, 2 and 4 iterations. Top left: the initial solution from the reaction rate equations immediately settles at one of the stable states. Top right and bottom left: the first and second switches are correctly obtained after respectively one and two iterations. Bottom right: all 4 spontaneous changes of state have been correctly detected. Note the rather large level of noise.

The relative errors and residuals are plotted in Figure 5.2. It is seen that the convergence initially is very rapid but then reaches a plateau where the error stays within a small fraction of the order of the noise level. The accuracy thus obtained is quite reasonable and is at the order of a 1% perturbation of the propensities in (5.3).

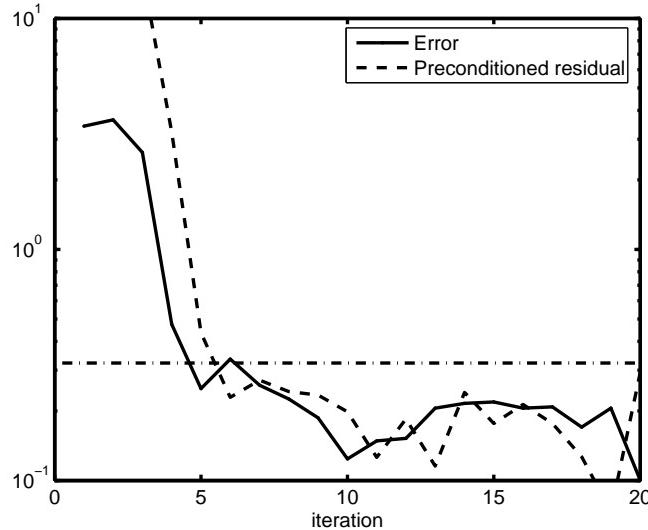


Figure 5.2: Maximum relative errors and residuals (see (5.1) and (5.2)) obtained for the toggle switch during the first 20 iterations. The accuracy is quickly improved during the first 5 iterations and then settles more slowly. The horizontal line (dash-dot) is the induced difference when the propensities are perturbed by $\pm 1\%$.

5.3 Homogenization of disparate rates

As a specific model containing two different scales we consider *fast dimerization* combined with *slow isomerization*,

$$\left. \begin{array}{l} X_1 + X_1 \xrightleftharpoons[1/\varepsilon]{1/\varepsilon} X_2 + X_2 \\ X_2 \xrightleftharpoons[1/\varepsilon]{1} Y_2 \\ Y_2 + Y_2 \xrightleftharpoons[1/\varepsilon]{1/\varepsilon} Y_1 + Y_1 \end{array} \right\}, \quad (5.4)$$

where the small parameter ε controls the difference in scales. For this example we took $\varepsilon = 10^{-3}$ and initial data $[x_1, x_2, y_1, y_2](0) = [15, 5, 30, 10]$ with final time $T = 10$. The fast channels thus fire about 10^4 times more often than the

slow ones and we note also that the fast channels have quadratic propensities (e.g. $w_1 = x_1(x_1 - 1)/\varepsilon$) so that the rate equations are inexact.

For the parareal discretization we used $\Delta t = T/50$ as before and an extremely simple coarse solver in the form of a single step with the linearized backward Euler method. The homogenization procedure described in Section 4.3 was used with the fine solver defined in (4.30) using $\delta t = \Delta t/2$. The resulting combination was very effective indeed in obtaining a homogenized solution, see Figure 5.3 and 5.4.

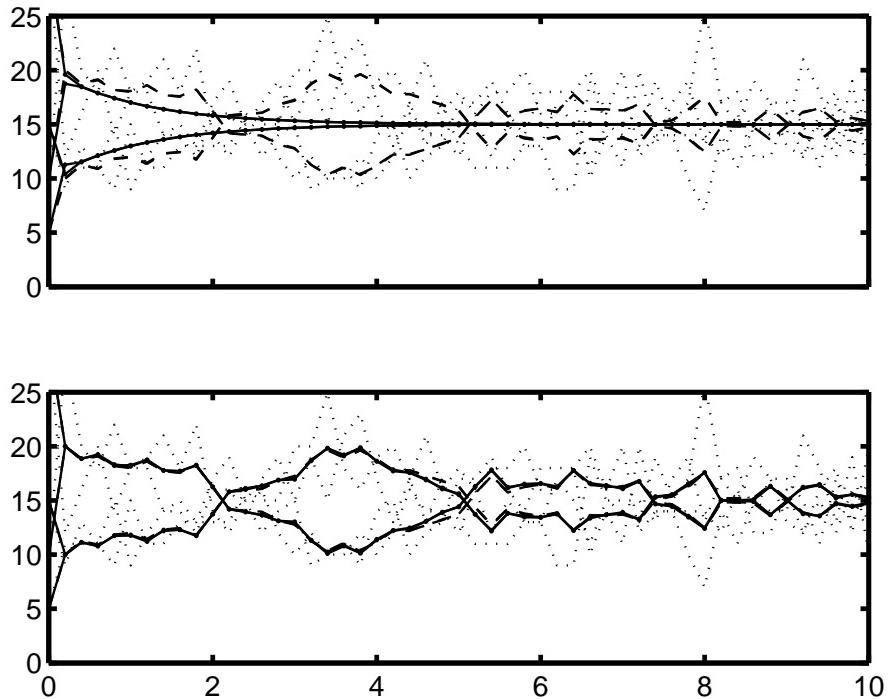


Figure 5.3: Dashed: homogenized solution for (5.4) versus time. Solid: solution after 0 (top) and 1 (bottom) iteration of the parareal algorithm. For comparison, the corresponding non-homogenized trajectory is also plotted using a dotted line. For this example, a single step of the backward Euler method was a sufficient accurate coarse scale solver which therefore is extremely cheap to evaluate.

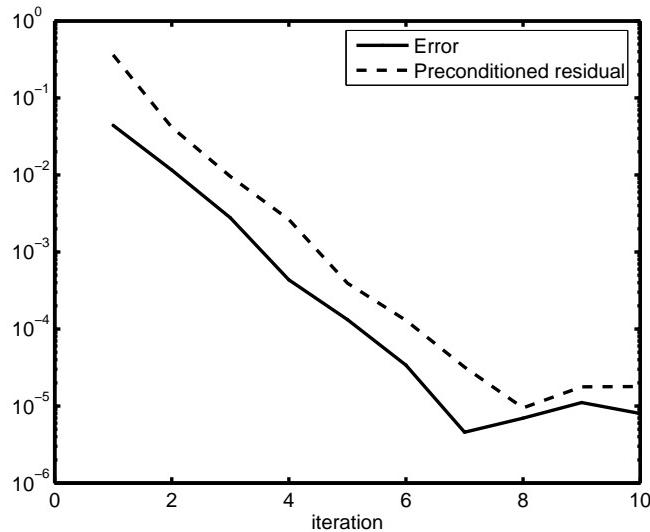


Figure 5.4: Convergence history for (5.4) homogenized according to (4.30).

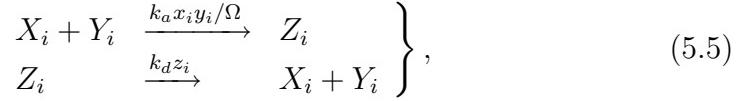
5.4 Stochastic reaction-diffusion

There are many interesting chemical systems where the spatial distribution of the species must be taken into account. The usual thermodynamical equilibria assumption (“well stirredness”) no longer holds as the transport of the molecules through the solvent is slow compared to average reaction times or when some reactions are strongly localized. The length scale over which the system can be regarded as homogeneous is now much shorter.

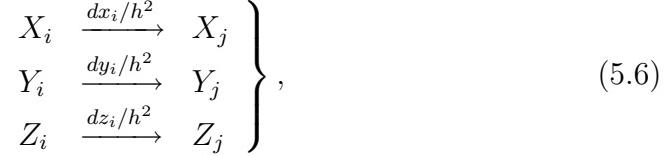
Examples of when both the stochasticity of the reactions and the spatial distribution are necessary to explain experimental data are found in [10, 12, 16, 30, 37]. For such systems, the diffusion at a molecular level can be treated as a special set of linear propensities. This yields the reaction-diffusion master equation (RDME) [18, Ch. 8], [29, Ch. XIV] which evolves the probability density of the system in the same manner as the CME (2.7). However, the dimensionality of the state-space is much higher and computing even a single trajectory can be a very computationally intensive problem. Note also that, as the rate equations now form the reaction-diffusion PDE, the master operator defines a continuous spectrum of scales so that a clear separation into slow/fast ones is no longer possible.

We shall consider a small example in the present section as follows. For $i = 1 \dots 5$, denote by the triplet $(x, y, z)_i$ the number of X -, Y - and Z -molecules in cell i . The cells each have volume Ω and for simplicity are

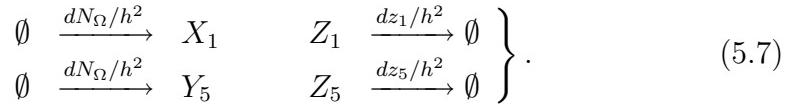
connected in an array 1-2-3-4-5. Within each cell i we specify the *reactions*



compare (2.2)–(2.5). Furthermore, the molecules can *diffuse* to any neighboring cell $j = i \pm 1 \in \{1, \dots, 5\}$ according to



where $h \equiv \Omega^{1/3}$ is the length-scale. Finally, there are also in- and outflow at the boundaries,



In order to fully describe the system we have chosen constants to approximately coincide with those in [12] for *E. coli*, see also [6]. The model is thus parameterized by the integer N_Ω defining the total volume through $V = 10^{-15} \times N_\Omega/25 = 5\Omega$ (l). The rate constants are given by $k_a = 10^8$ ($M^{-1}s^{-1}$), $k_d = 10$ (s^{-1}) and $d = 10^{-10}$ (m^2s^{-1}). Since we have that $N_\Omega \propto \Omega$, for convenience, we loosely refer to both quantities as the “system size”. The model is normalized around $N_\Omega = 25$ molecules per cell and as initial data we simply took $(x_i, y_i, z_i) = (N_\Omega, N_\Omega, 0)$ in all cells.

For the parareal algorithm we again took $N = 50$ intervals with total time $T = 1$. In the semi-discrete case, one can show that the reaction-rate equations for the diffusion part (5.6) are equivalent to a mass-lumped FEM-method for the macroscopic diffusion equation; this observation has been used to generalize the mesoscopic model to more complicated geometries [14]. We note in passing that this setting opens up for extremely efficient coarse-grain solvers based on multigrid techniques when the spatial resolution increases.

In Figure 5.5 the error during each parareal step is displayed. There is a trend with faster convergence as the system size increases and despite the seemingly rather large error levels, the solution obtained is visually very pleasing and can hardly be distinguished from the exact one (cf. Figure 5.7). In Figure 5.6 we have used instead the homogenized fine scale solver (4.30) and the convergence improves a lot. Again, smoothing the output from the fine solver by integrating over a short period of time has the effect of scaling down the effective Lipschitz constant and convergence to the homogenized solution becomes fast.

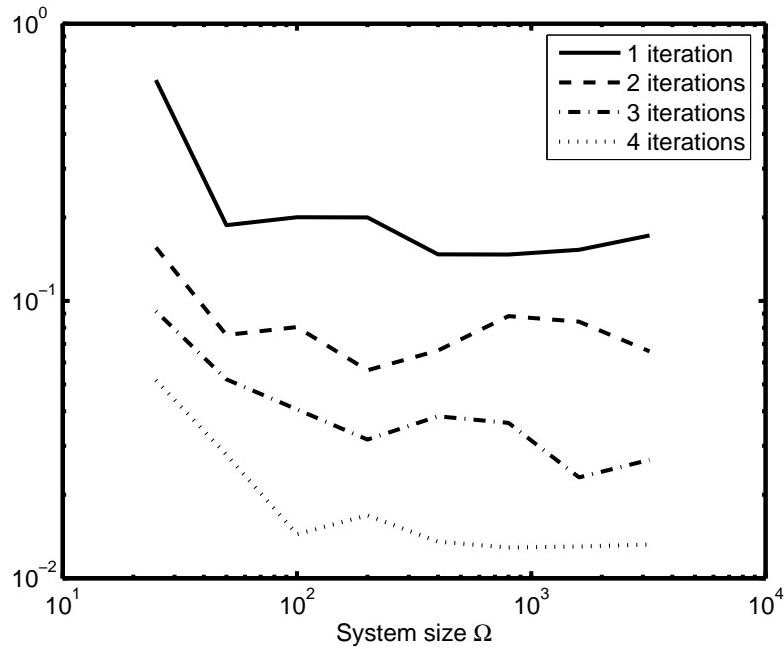


Figure 5.5: Relative error per the first few parareal iterations when the system size Ω increases.

6 Conclusions

The master equation, or equivalently, a special type of jump SDEs accurately describes well stirred chemical reactions at the mesoscopic level. Diffusion can be incorporated by a special set of reactions implying that non-homogenous problems also can be modeled using the same type of process. In the macroscopic limit the usual rate equations (or the reaction-diffusion PDE) emerges.

The parareal algorithm can be applied by using the rate equations as a predictor for the jump process. Convergence is then dictated by the Lipschitz constant of the system, but also by the level of noise in the solution. This noise vanishes in the macroscopic limit.

For models involving rapid transients the best one can generally hope for is convergence in moment, but it is then not clear how the error should be monitored. A remedy is to homogenize the model on the fly by averaging the process in time. This homogenization is very general and should be applicable to a wide range of stochastic time-dependent problems. We remark that obtaining this homogenized solution is impractical on a serial computer since the full solution must be obtained first. Furthermore, the proposed

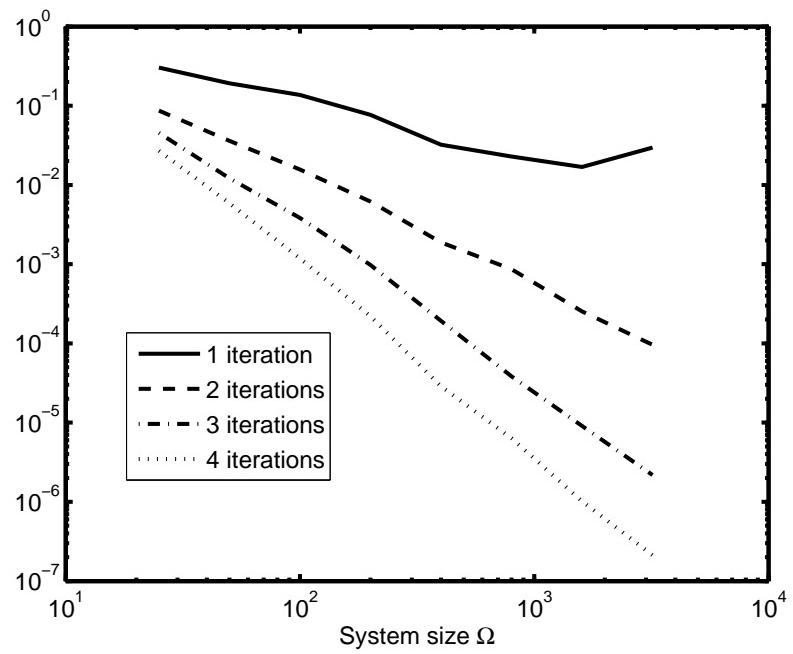


Figure 5.6: As in Figure 5.5 but using a homogenized fine solver (4.30) with $\delta t = \Delta t / 4$. The convergence to the homogenized solution is faster and more regular.

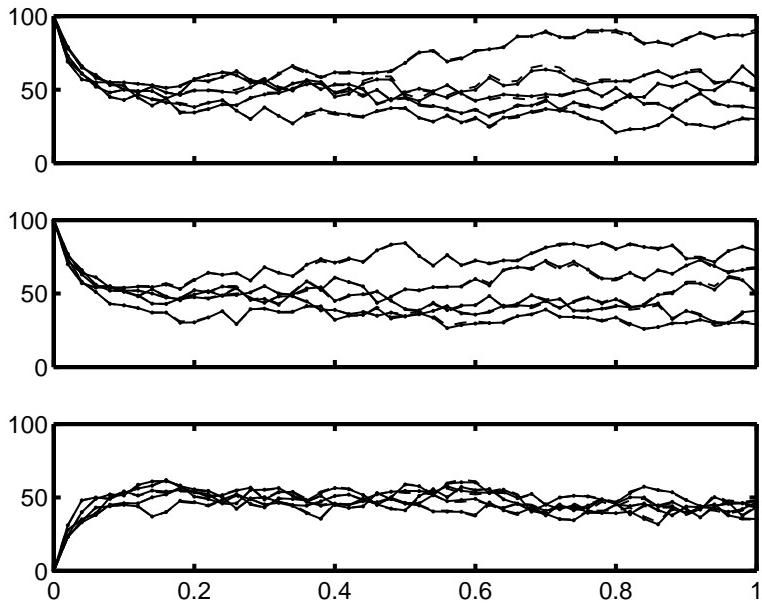


Figure 5.7: Sample solution for the reaction-diffusion system (5.5), (5.6) and (5.7). The system-size parameter is for this example $N_\Omega = 100$ molecules per cell. Solid: solution obtained after 4 parareal iterations over 50 processors, dashed and barely visible: exact trajectory. From top to bottom are the trajectories for the X -, the Y - and the Z -molecules, respectively. Towards the end of the simulation there is a strong negative correlation between X - and Y -molecules and there is also a spatial correlation due to the unsymmetric inflow (5.7).

homogenization can be put in contrast to other multiscale methods [7, 11]. Typically one computes a solution to a homogenized *equation* obtained by determining some kind of averaged coefficients. With the parareal algorithm, one rather directly obtains a homogenized *solution*.

The range of applicability of the method can be increased by adding a white-noise term to the governing equations, thus encompassing highly general Lévy-type SDEs [2]. In this general setting, the suggested homogenization provides a way of applying the parareal algorithm to stiff problems.

We would also like to relate our results to those in reference [5] mentioned in the introduction. There it is shown that the parareal/Euler forward combination for a *nonstiff* Wiener SDE yields a mean square estimate $\mathcal{O}(\Delta t^{k/2})$ for the k th parareal iteration assuming a practically exact fine solver. This holds true in the present context as well provided that we use the appropriate version of the forward Euler method, namely the tau-leap method [31]. *However*, the resulting method is highly nontrivial, if possible at all, to implement efficiently. The reason is that the same realizations of Poisson processes have to be used for both the coarse and the fine solver. The list of events simulated by the fine solver must somehow be reused and searched through leading to a very expensive coarse solver. An open question is thus if one can somehow circumvent this issue.

One of the most promising applications of the proposed method seems to be for reaction-diffusion models such as the one investigated in Section 5.4. Here the macroscopic model can typically be acceptable in many, but not all, subvolumes. Some of the species are typically present in fairly large copy numbers where noise is less pronounced. A few iterations of the suggested parareal algorithm can thus be understood as a kind of deterministic/stochastic hybrid method. A feature with this set-up is that one never needs to explicitly determine what parts should be treated deterministically.

Acknowledgment

Part of this work was initiated during a 7 month visit to the Division of Applied Mathematics at Brown University, Providence, Rhode Island. During this visit Jan Hesthaven and his family kindly provided with both practical and personal support. At Brown University, Yvon Maday supplied several references on the parareal algorithm.

During the time of writing, Ingemar Kaj offered valuable initial suggestions and a much appreciated discussion. Finally, Per Lötstedt and Andreas Hellander made several comments that improved the quality of the manuscript.

References

- [1] D. F. Anderson. Incorporating postleap checks in tau-leaping. *J. Chem. Phys.*, 128(5):054103–054111, 2008. doi:10.1063/1.2819665.
- [2] D. Applebaum. *Lévy Processes and Stochastic Calculus*, volume 93 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2004.
- [3] L. Baffico, S. Bernard, Y. Maday, G. Turinici, and G. Zérah. Parallel-in-time molecular-dynamics simulations. *Phys. Rev. E*, 66(5):057701–057704, 2002. doi:10.1103/PhysRevE.66.057701.
- [4] G. Bal. On the convergence and the stability of the parareal algorithm to solve partial differential equations. In *Domain Decomposition Methods in Science and Engineering*, volume 40 of *Lecture Notes in Computational Science and Engineering*, pages 425–432. Springer, Berlin, 2005. doi:10.1007/3-540-26825-1_43.
- [5] G. Bal. Parallelization in time of (stochastic) ordinary differential equations, 2006. Submitted. Available at <http://www.columbia.edu/~gb2030>.
- [6] O. G. Berg and P. H. von Hippel. Diffusion-controlled macromolecular interactions. *Ann. Rev. Biophys. Chem.*, 14(1):131–160, 1985. doi:10.1146/annurev.bb.14.060185.001023.
- [7] Y. Cao, D. Gillespie, and L. Petzold. Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems. *J. Comput. Phys.*, 206:395–411, 2005. doi:10.1016/j.jcp.2004.12.014.
- [8] Y. Cao, L. R. Petzold, M. Rathinam, and D. T. Gillespie. The numerical stability of leaping methods for stochastic simulation of chemically reacting systems. *J. Chem. Phys.*, 121(24):12169–12178, 2004. doi:10.1063/1.1823412.
- [9] D. R. Cox and V. Isham. *Point processes*. Monographs on applied probability and statistics. Chapman and Hall, London, 1980.
- [10] M. Dobrzański, J. V. Rodríguez, J. A. Kaandorp, and J. G. Blom. Computational methods for diffusion-influenced biochemical reactions. *Bioinformatics*, 23(15):1969–1977, 2007. doi:10.1093/bioinformatics/btm278.

- [11] W. E, D. Liu, and E. Vanden-Eijnden. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *J. Chem. Phys.*, 123(19), 2005. doi:10.1063/1.2109987.
- [12] J. Elf and M. Ehrenberg. Spontaneous separation of bi-stable biochemical systems into spatial domains of opposite phases. *Syst. Biol.*, 1(2):230–236, 2004. doi:10.1049/sb:20045021.
- [13] M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002. doi:10.1126/science.1070919.
- [14] S. Engblom, L. Ferm, A. Hellander, and P. Lötstedt. Simulation of stochastic reaction-diffusion processes on unstructured meshes. Technical Report 2008-012, Dept of Information Technology, Uppsala University, Uppsala, Sweden, 2008. Available at <http://www.it.uu.se/research>.
- [15] S. N. Ethier and T. G. Kurtz. *Markov processes: characterization and convergence*. Wiley series in Probability and Mathematical Statistics. John Wiley & Sons, New York, 1986.
- [16] D. Fange and J. Elf. Noise-induced min phenotypes in *E. coli*. *PLoS Comput. Biol.*, 2(6):637–648, 2006. doi:10.1371/journal.pcbi.0020080.
- [17] M. J. Gander and S. Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM J. Sci. Comput.*, 29(2):556–578, 2007. doi:10.1137/05064607X.
- [18] C. W. Gardiner. *Handbook of Stochastic Methods*. Springer Series in Synergetics. Springer-Verlag, Berlin, 3rd edition, 2004.
- [19] T. S. Gardner, C. R. Cantor, and J. J. Collins. Construction of a genetic toggle switch in Escherichia coli. *Nature*, 403:339–342, 2000. doi:10.1038/35002131.
- [20] M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem.*, 104(9):1876–1889, 2000. doi:10.1021/jp993732q.
- [21] I. I. Gihman and A. V. Skorohod. *Stochastic differential equations*, volume 72 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, Berlin, 1972.

- [22] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.*, 22(4):403–434, 1976. doi:[10.1016/0021-9991\(76\)90041-3](https://doi.org/10.1016/0021-9991(76)90041-3).
- [23] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188:404–425, 1992. doi:[10.1016/0378-4371\(92\)90283-V](https://doi.org/10.1016/0378-4371(92)90283-V).
- [24] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.*, 115(4):1716–1733, 2001. doi:[10.1063/1.1378322](https://doi.org/10.1063/1.1378322).
- [25] P. Guptasarama. Does replication-induced transcription regulate synthesis of the myriad low copy number proteins of *Escherichia coli*? *Bioessays*, 17(11):987–997, 1995. doi:[10.1002/bies.950171112](https://doi.org/10.1002/bies.950171112).
- [26] E. L. Haseltine and J. B. Rawlings. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys.*, 117(15), 2002. doi:[10.1063/1.1505860](https://doi.org/10.1063/1.1505860).
- [27] N. Ikeda and S. Watanabe. *Stochastic Differential Equations and Diffusion Processes*, volume 24 of *North-Hollans Mathematical Library*. North-Holland/Kodansha, Amsterdan/Tokyo, 1981.
- [28] J. Jacod and A. N. Shiryaev. *Limit Theorems for Stochastic Processes*, volume 288 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, 1987.
- [29] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, Amsterdam, 5th edition, 2004.
- [30] K. Kruse and J. Elf. Kinetics in spatially extended systems. In Z. Szallasi, J. Stelling, and V. Periwal, editors, *System Modeling in Cellular Biology. From Concepts to Nuts and Bolts*, pages 177–198. MIT Press, Cambridge, MA, 2006.
- [31] T. Li. Analysis of explicit tau-leaping schemes for simulating chemically reacting systems. *Multiscale Model. Simul.*, 6(2):417–436, 2007. doi:[10.1137/06066792X](https://doi.org/10.1137/06066792X).
- [32] T. Li, A. Abdulle, and W. E. Effectiveness of implicit methods for stiff stochastic differential equations. *Commun. Comput. Phys.*, 3(2):295–307, 2008.

- [33] J.-L. Lions, Y. Maday, and G. Turinici. Résolution d'EDP par un schéma en temps “pararéel”. *C. R. Acad. Sci. Paris Sér. I Math.*, 332(7):661–668, 2000. doi:10.1016/S0764-4442(00)01793-6.
- [34] Y. Maday. Parareal in time algorithm for kinetic systems based on model reduction. In A. Bandrauk, M. C. Delfour, and C. Le Bris, editors, *High-Dimensional Partial Differential Equations in Science and Engineering*, volume 41 of *CRM Proceedings & Lecture Notes*, pages 183–194. AMS/CRM, 2007.
- [35] Y. Maday and G. Turinici. A parareal in time procedure for the control of partial differential equations. *C. R. Acad. Sci. Paris Sér. I Math.*, 335(4):387–392, 2002. doi:10.1016/S1631-073X(02)02467-6.
- [36] H. H. McAdams and A. Arkin. It’s a noisy business. genetic regulation at the nanomolar scale. *Trends Gen.*, 15(2):65–69, 1999. doi:10.1016/S0168-9525(98)01659-X.
- [37] R. Metzler. The future is noisy: The role of spatial fluctuations in genetic switching. *Phys. Rev. Lett.*, 87(6):068103–068107, 2001. doi:10.1103/PhysRevLett.87.068103.
- [38] J. Paulsson, O. G. Berg, and M. Ehrenberg. Stochastic focusing: Fluctuation-enhanced sensitivity of intracellular regulation. *Proc. Natl. Acad. Sci. USA*, 97(13):7148–7153, 2000. doi:10.1073/pnas.110057697.
- [39] S. Plyasunov. On hybrid simulation schemes for stochastic reaction dynamics, 2005. Available at <http://arxiv.org/abs/math/0504477>.
- [40] J. M. Raser and E. K. O’Shea. Noise in gene expression: Origins, consequences and control. *Science*, 309(5743):2010–2013, 2005. doi:10.1126/science.1105891.
- [41] M. Rathinam, L. Petzold, Y. Cao, and D. T. Gillespie. Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *J. Chem. Phys.*, 119(24):12784–12794, 2003. doi:10.1063/1.1627296.
- [42] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie. Consistency and stability of tau-leaping schemes for chemical reaction systems. *Multiscale Model. Simul.*, 4(3):867–895, 2005. doi:10.1137/040603206.

- [43] G. A. Staff. The parareal algorithm, 2003. Available at www.idi.ntnu.no/~elster/notur-cluster03.
- [44] G. A. Staff, Y. Maday, and E. M. Rønquist. The parareal-in-time algorithm: basics, stability analysis and more, 2006. Submitted. Available at <http://www.springerlink.com/index/1p436711n8174312.pdf>.
- [45] Y. Togashi and K. Kaneko. Molecular discreteness in reaction-diffusion systems yields steady states not seen in the continuum limit. *Phys. Rev. E*, 70(2):020901–1–020901–4, 2004. doi:[10.1103/PhysRevE.70.020901](https://doi.org/10.1103/PhysRevE.70.020901).
- [46] J. M. G. Vilar, H. Y. Kueh, N. Barkai, and S. Leibler. Mechanism of noise-resistance in genetic oscillators. *Proc. Nat. Acad. Sci.*, 99:5988–5992, 2002. doi:[10.1073/pnas.092133899](https://doi.org/10.1073/pnas.092133899).